

DIFFUSION LIMITS FOR NETWORKS OF MARKOV-MODULATED INFINITE-SERVER QUEUES

H. M. JANSEN^{1,2}, M. MANDJES¹, K. DE TURCK³, S. WITTEVRONGEL²

ABSTRACT. This paper studies the diffusion limit for a network of infinite-server queues operating under Markov modulation (meaning that the system's parameters depend on an autonomously evolving background process). In previous papers on (primarily single-node) queues with Markov modulation, two variants were distinguished: one in which the server speed is modulated, and one in which the service requirement is modulated (i.e., depends on the state of the background process upon arrival). The setup of the present paper, however, is more general, as we allow both the server speed and the service requirement to depend on the background process. For this model we derive a Functional Central Limit Theorem: we show that, after accelerating the arrival processes and the background process, a centered and normalized version of the network population vector converges to a multivariate Ornstein-Uhlenbeck process. The proof of this result relies on expressing the queueing process in terms of Poisson processes with a random time change, an application of the Martingale Central Limit Theorem, and continuous-mapping arguments.

KEYWORDS. Queueing networks ◦ infinite-server queue ◦ diffusion limit ◦ Markov modulation

¹ Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, the Netherlands.

² TELIN, Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium.

³ CentraleSupélec, Département de Télécommunications, Laboratory of Signals and Systems (L2S), UMR8506 Plateau de Moulon, 3 rue Joliot-Curie, 91192 Gif sur Yvette, France.

E-mail. {h.m.jansen|m.r.h.mandjes}@uva.nl, koen.deturck@supelec.fr, sw@telin.ugent.be

M. Mandjes' research is partly funded by the NWO Gravitation Programme NETWORKS, Grant Number 024.002.003, and an NWO Top Grant, Grant Number 613.001.352.

1. INTRODUCTION

Networks of infinite-server queues with time-dependent parameters have been extensively studied. In a series of articles, several nice properties have been shown for the situation that the network's queues are fed by Poisson arrival streams with time-dependent deterministic arrival rates. In this model, due to the fact that customer streams do not interact, departure streams are again time-dependent Poisson (under specific assumptions on the queues' initial conditions). In addition, at any point in time the joint queue length has independent Poisson marginals. Without pursuing an exhaustive overview, we refer to the contributions by Harrison and Lemoine [13] and Eick *et al.* [9].

Modelling customer streams as Poisson processes with time-dependent deterministic arrival rates has, however, the following crucial drawback. Due to standard properties of such processes, for any time interval the mean number of arrivals coincides with the corresponding variance. In various practical contexts, however, there is *overdispersion*, i.e., the variance substantially exceeds the mean; see e.g. [3, 17] and references therein. To incorporate overdispersion a common procedure is to apply *modulation*, meaning that the arrival rate depends on an independently evolving stochastic process (usually referred to as the background process). Due to the additional uncertainty that is introduced by the background process, the variance of the number of arrivals now exceeds the

mean. An example of modulation is to let the arrival process be a Cox process, i.e., a Poisson process whose rate is a (nonnegative) stochastic process.

Above we introduced the concept of modulation for arrival processes, but it can clearly be applied to the queues' service processes as well. It is important to notice that modulation of service processes comes in multiple flavors. In the first place, we can modulate the server speed. A slightly more complicated form of modulation lets the service requirement distribution depend on the background process: the service requirement is sampled upon arrival, with a parameter determined by the state of the background process at the moment the customer arrives. Thus, based on which parameters of the model are modulated, the following four types of modulation may be distinguished. In line with the terminology used in [5],

- Model 0: only the arrival rate is modulated,
- Model I: only the arrival rate and the server speed are modulated,
- Model II: only the arrival rate and the service requirement distribution are modulated,
- Model III: all parameters are modulated.

Obviously, Model 0 is a special case of both Model I and Model II, and Model I and Model II are special cases of Model III.

Introducing modulation complicates the analysis of networks of infinite-server queues considerably: many of the nice properties that were found for time-dependent deterministic arrival rates do not carry over to the setting in which the arrival rate is stochastic, for instance. More specifically, the joint queue length ceases to have independent marginals, and in addition these marginals are not Poisson. An important strand of work is focused on *Markov modulation*: the value of the model parameters is determined by an external finite-state continuous-time Markov chain. Early work on (primarily single-node) Markov-modulated infinite-server systems, such as [19], develops procedures to identify the (transient and stationary) moments of the number of clients; see also [4, 8, 12] for more recent contributions. A significant drawback of most of these results concerns their implicit nature: systems of partial differential equations are derived for the moment generating function pertaining to the joint distribution of the number of clients in the individual queues, and recursions for the corresponding moments are found. Only in very special cases the underlying probability distribution can be found in more explicit terms; see e.g. [11].

The lack of explicit expressions motivated the systematic study of scaling regimes in which the queue length distribution *can* be given in closed form. Some recent papers have established scaling limits in the context of single-queue Model 0 systems: [14] considers the scenario in which the arrival rate is periodically resampled, whereas [18] analyzes the case in which the arrival rate is given by a shot-noise process. For the case of single queues under Markov modulation, it is shown in [2] (for Model 0) and [5] (for Model I and Model II) that, after appropriately accelerating the arrival processes and the background process, a centered and normalized version of the number of customers in the system converges to an Ornstein-Uhlenbeck process. The proof of this Functional Central Limit Theorem (FCLT) relies on casting the queueing process in terms of Poisson processes with a random time change and an application of the Martingale Central Limit Theorem.

The main goal of the present paper is to extend the findings of [2, 5] in two directions. First of all, we study a network of infinite-server queues under Markov modulation (rather than a single queue). Next to that, we do not restrict ourselves to Model 0, Model I, or Model II, but focus on the Model III variant which covers the other variants as special cases. Our contribution is that we derive, for this highly general setting, an FCLT.

The different setting that we study also motivates the use of a different strategy to establish an FCLT. As will become clear in the paper, the argumentation for this multivariate setting has some elements in common with the one used in the single-queue case, but in general requires considerably more delicate arguments. A key insight is that we may exploit continuous-mapping arguments once we have established asymptotic equivalence of our network with some other, less complicated stochastic process. Following this strategy, we succeed in establishing an FCLT: the centered and normalized version of the joint queue length process converges to a multivariate Ornstein-Uhlenbeck process.

The paper is organized as follows. In Section 2, we introduce our model and point out how to express the network population vector in terms of unit-rate Poisson processes with a random time change. We also show that, perhaps somewhat surprisingly, any network of Model III type can be expressed in terms of a, seemingly more restrictive, Model I network, entailing that it suffices to just consider the Model I setting from then on. Then Section 3 presents the fluid limit under a scaling of both the arrival rates and the background process. Centering the queue lengths with this fluid limit, the results in Section 4 show that, after a suitable normalization, the joint queue length process weakly converges to a multivariate Ornstein-Uhlenbeck process. Proofs and auxiliary results are given in Section 5.

2. MODEL AND PRELIMINARIES

In this section we define a Markov-modulated network of infinite-server queues with exponentially distributed service requirements. We introduce notation for the Model I variant and show how the queue length vector can be represented in terms of Poisson processes with a random time change. We then argue that, in the setting that we consider, a network containing Model II or Model III queues can be represented as a network containing only Model I queues (justifying that we restrict ourselves to Model I in the rest of the paper).

2.1. Definition of Model I. We now provide the mathematical description of the Model I version of a Markov-modulated network of infinite-server queues. We assume that the background process J is an autonomously evolving continuous-time Markov chain with finite state space $\{1, \dots, d\}$ and irreducible transition rate matrix \mathcal{Q} . The invariant distribution corresponding to \mathcal{Q} is denoted by the $d \times 1$ column vector π (so π has nonnegative entries that sum to 1 and $\pi^\top \mathcal{Q} = 0$).

Jobs that leave a particular queue may be routed to another queue or leave the network altogether. Throughout the paper we assume there are $L \in \{2, 3, \dots\}$ queues in the network. We impose the assumption of at least two queues, because this allows us to not explicitly record the jobs leaving the system. Indeed, a job leaving the system is equivalent to sending it to some designated infinite-server queue that has zero server speed (i.e., a queue from which jobs cannot leave). The assumption $L \in \{2, 3, \dots\}$ is therefore not a restriction and is merely intended to streamline notation.

Given a network of L queues, we define the set of queue indices $\mathcal{I} = \{1, \dots, L\}$. For any $k \in \mathcal{I}$, we define $\mathcal{I}_k = \mathcal{I} \setminus \{k\}$. This is, of course, the set of all queues to which a job from queue k may be routed after service completion. This set is always nonempty, because $L \in \{2, 3, \dots\}$ by assumption.

When the background process is in $i \in \{1, \dots, d\}$, jobs arrive at queue k as a Poisson process with rate $\lambda_k(i)$. In addition, again for the background process being in state i , we denote by $\mu_{kl}(i)$ the rate at which any individual job present in queue k completes service and jumps to queue l . Observe that during the time a job spends in a given queue, the rate of jumping to another queue may change in time (when the state of the background process changes, that is).

We denote the number of jobs in queue $k \in \mathcal{I}$ at time $t \geq 0$ by $Q_k(t)$, and the L -dimensional queue content process by

$$Q(t) = (Q_1(t), \dots, Q_L(t)). \quad (1)$$

Then we represent Q as the solution of the system of equations

$$\begin{aligned} Q_k(t) = & Q_k(0) + A_k \left(\int_0^t \lambda_k(J(s)) ds \right) + \sum_{l \in \mathcal{I}_k} S_{lk} \left(\int_0^t \mu_{lk}(J(s)) Q_l(s) ds \right) \\ & - \sum_{l \in \mathcal{I}_k} S_{kl} \left(\int_0^t \mu_{kl}(J(s)) Q_k(s) ds \right). \end{aligned} \quad (2)$$

Here, $Q(0)$ is an independent random vector taking values in $\mathbb{Z}_{\geq 0}^L$, with its k -th entry describing the initial number of jobs in queue k . The processes A_k and S_{kl} with $k \in \mathcal{I}$ and $l \in \mathcal{I}_k$ form a collection of independent standard (unit-rate, that is) Poisson processes. We use A_k to model the exogenous arrivals to queue k , whereas we use S_{kl} to model service completions in queue k that are routed to queue l . This representation has been extensively used in e.g. [20], and is (in our experience at least) an intuitive representation that is relatively easy to work with when studying weak convergence.

2.2. Models II and III are covered by Model I. In the network setting that we consider, it turns out, perhaps somewhat surprisingly, that any Model II or Model III network of infinite-server queues can be described through a Model I network of infinite-server queues. At first sight, our construction of a network of Model I queues may not seem very flexible, because the modulation of service requirements (in Model II and Model III queues) is not explicitly incorporated. However, the combination of Model I queues with the network setting gives us much more flexibility than we might have thought.

The upcoming examples are intended to demonstrate this flexibility by explicitly showing how we may incorporate modulated service requirements. In particular, the second example demonstrates how we may construct a Model III queue as a network of Model I queues, thus showing that a network of Model I queues is indeed the most general Markov-modulated network of infinite-server queues.

Remark 2.1. Observe that the Model I setting incorporates probabilistic routing: when the background process is in state i , the rate of leaving queue k is $\mu_k(i) := \sum_{l \neq k} \mu_{lk}(i)$, and the probability that after being served the job jumps to queue l is $p_{kl}(i) := \mu_{kl}(i)/\mu_k(i)$. As such, the routing probabilities depend on the state of the background process at service completion. \diamond

Example 2.2. In Remark 2.1 we already argued that Model I covers settings in which the routing probabilities depend on the current state of the background process. An obvious extension would be to make the routing probabilities depend on the state of the background process at the moment of a job's arrival. In this example we argue that such a mechanism still falls in the context of Model I networks of Markov-modulated infinite-server queues.

We consider a network in which jobs arrive to queue A at rate $\lambda_A(J(t))$ and receive service there with speed $\mu_A(J(t))$. When a job completes service at this queue, it is routed with probability $p(i, j)$ to queue B, where i is the state of the background process J upon arrival of the job and j is the state of J upon service completion. With probability $1 - p(i, j)$ the job is routed to queue C. (What happens at queue B and C is not very relevant here; neither do we need a queue D for jobs leaving the system.) The difficulty here is the dependence of the routing probabilities on both the present and the past. Indeed, the routing probabilities for a job are determined when it completes service, but also depend on the state of J upon its arrival.

We can deal with this problem by registering the state of the background process J when a job arrives to queue A. This naturally leads to the jobs in queue A being divided into d different classes (where we recall that d is the size of the state space of J): a job belongs to class i when J is in state i upon arrival of the job. The idea is now to split the content of queue A into d individual queues, i.e., one for each class. Then the routing probabilities for jobs in those queues only depend on the queue index and the state of J upon service completion, which is in line with the probabilistic routing mechanism that we discussed in Remark 2.1.

In more concrete terms, the system considered as a network of Model I queues will look like this. We denote the queues handling the different classes by $1, \dots, d$. We define the arrival rates to these queues by $\lambda_k(i) = \mathbb{1}_{\{i=k\}}\lambda_A(i)$ for $k, i \in \{1, \dots, d\}$. Then we have arrivals of only one class at a time with the right arrival rate for each class.

Jobs from those d queues will be routed to queue $d+1$ and queue $d+2$ (which were queue B respectively queue C in our original thought experiment). The routing mechanism is then incorporated by defining the server speeds by picking $\mu_{k(d+1)}(i) = p(k, i)\mu_A(i)$ and $\mu_{k(d+2)}(i) = (1 - p(k, i))\mu_A(i)$ for $k \in \{1, \dots, d\}$. \diamond

Example 2.3. In the previous example, we have seen how we can handle different classes (corresponding to the state of the background process that the job sees upon arrival) by assigning them to their own Model I queue. In this example, we use this idea to show how a Model III Markov-modulated infinite-server queue may be represented as a Model I Markov-modulated network of infinite-server queues.

Consider a single Model III Markov-modulated infinite-server queue where J is the background process. In this case, jobs arrive according to a Poisson process with rate $\lambda^*(i)$ while J is in state i . If a job arrives to the system while J is in state i , then it will be of type i and has an independent service requirement having an exponential distribution with parameter $\kappa^*(i)$. The server speed is $\mu^*(i)$ while the background process is in state i .

We will construct a network of $L = d + 1$ queues: d queues for each job type (i.e., queue k contains the jobs of type k that are still in service) and one queue to collect all jobs that have finished service in their respective queue (i.e., queue $d + 1$ with zero exogenous arrival rate and zero server speed). In what follows, i ranges over $\{1, \dots, d\}$, and k and l range over $\{1, \dots, d, d + 1\}$.

To construct the network, define $\lambda_k(i) = \mathbb{1}_{\{i=k\}}\lambda^*(i)$ for $k \in \{1, \dots, d\}$ and $\lambda_{d+1}(i) = 0$. This means that jobs arrive to exactly one queue at a time: the arrival rate (of jobs of type i) to queue i is $\lambda^*(i)$ while J is in state i and zero while J is not in state i . Moreover, there are no exogenous arrivals to queue $d + 1$, which collects jobs that have finished service. Additionally, we take $\mu_{kl}(i) = \mathbb{1}_{\{l=d+1\}}\kappa^*(k)\mu^*(i)$ for $k \in \{1, \dots, d\}$. This ensures that jobs can only be routed to queue $d + 1$ after service completion (the $\mathbb{1}_{\{l=d+1\}}$ part), that jobs of type k arriving at queue k have an exponentially distributed service requirement with parameter $\kappa^*(k)$, and that those jobs in queue k experience server speed $\mu^*(i)$ while J is in state i .

We have thus defined a Model I network of $d + 1$ infinite-server queues, $Q_1(t)$ up to $Q_{d+1}(t)$, that is equivalent to our single Model III Markov-modulated infinite-server queue. Now the queue content process of this single Model III queue at time is given by $\sum_{k=1}^d Q_k(t)$. As a bonus, the departure process of the Model III queue is captured by $Q_{d+1}(t)$. Thus, the Model I network constructed above completely describes the behavior of the queue content of the Model III queue.

Along the same lines, it can be argued that a network of L Model III Markov-modulated infinite-server queues can be described as a network of Ld Model I Markov-modulated infinite-server queues. We conclude that we can restrict ourselves to Model I. \diamond

Remark 2.4. It is directly seen that phase-type service requirements are also easy to deal with, by introducing infinite-server queues corresponding to each of the phases. \diamond

3. CONVERGENCE RESULTS: A FLUID LIMIT

As mentioned earlier, we are interested in proving a diffusion limit for our Markov-modulated network of infinite-server queues. A first step towards such a result is the derivation of a fluid limit for the network. This serves three purposes. First, the fluid limit is interesting in its own right, as it describes the typical behavior of the network under scaling. It is, essentially, a law of large numbers for the evolution of the network population. It holds under more general assumptions than the diffusion limit. Second, in the proof for the diffusion limit we will apply the fluid limit to identify the limiting process; the fluid limit limit is also used to appropriately center the network population process. Third, we will use the fluid limit to point out differences between the modulated network and its nonmodulated counterpart.

To establish the fluid limit result, we scale the parameters of the modulated network. We do so by scaling the arrival rates to the queues, the server speeds, and the transition rates of the background process. In addition, the network's initial population is scaled appropriately. More specifically, we impose the following assumption.

Assumption 3.1. *In the n -th system,*

- with $Q_k^n(0)$ denoting the initial number of jobs in queue k , we assume that $\frac{1}{n}Q^n(0)$ converges in probability to the constant $\rho(0)$;
- with λ_k^n the arrival rate to queue k , we assume that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \lambda_k^n = \lambda_k; \quad (3)$$

- with μ_{kl}^n the service speed for jobs to be routed from queue k to queue l , we assume that

$$\lim_{n \rightarrow \infty} \mu_{kl}^n = \mu_{kl}; \quad (4)$$

- the background process, denoted by J^n , corresponds to a continuous-time Markov chain having irreducible generator matrix $n^\alpha \mathcal{Q}$ (for some $\alpha > 0$) and invariant distribution π .

Regarding the μ_{kl}^n and λ_k^n , the main examples to keep in mind are $\lambda_k^n = n\lambda_k$ and $\mu_{kl}^n = \mu_{kl}$.

Recalling that $J^n(t)$ matches with $J(n^\alpha t)$, we observe that the scaling proposed here amounts to simultaneously speeding up the arrivals and the time scale of the background process, while the server speeds approach some limiting value. Note that the speedup of the arrival rate is essentially linear in n , whereas the speedup of the time scale of the background process is sublinear ($\alpha < 1$), linear ($\alpha = 1$), or superlinear ($\alpha > 1$) in n .

Under this scaling, the queue content process is denoted by the \mathbb{R}^L -valued stochastic process Q^n , which is given by a system of equations involving unit-rate Poisson processes with a random time change: for $k = 1, \dots, L$ we have

$$\begin{aligned} Q_k^n(t) = & Q_k^n(0) + A_k \left(\int_0^t \lambda_k^n(J^n(s)) ds \right) + \sum_{l \in \mathcal{I}_k} S_{lk} \left(\int_0^t \mu_{lk}^n(J^n(s)) Q_l^n(s) ds \right) \\ & - \sum_{l \in \mathcal{I}_k} S_{kl} \left(\int_0^t \mu_{kl}^n(J^n(s)) Q_k^n(s) ds \right). \end{aligned} \quad (5)$$

(See also Eq. 2.) As mentioned, we are interested in finding a fluid limit for Q^n . This will involve the time-averaged parameters $\lambda_k^\pi = \sum_{i=1}^d \pi(i) \lambda_k(i)$ and $\mu_{kl}^\pi = \sum_{i=1}^d \pi(i) \mu_{kl}(i)$. We present the fluid limit in the next lemma.

Lemma 3.2. *Under Assumption 3.1, the process $\frac{1}{n}Q^n$ converges uoc in probability to the solution of the system of integral equations*

$$\rho_k(t) = \rho_k(0) + \int_0^t \lambda_k^\pi ds + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}^\pi \rho_l(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}^\pi \rho_k(s) ds, \quad (6)$$

for $k = 1, \dots, L$.

This fluid limit is exactly the fluid limit that we would obtain from a *nonmodulated* network of infinite-server queues, and is in particular not affected by the value of α . (To see this, just replace the modulated parameters of the queues by their time-averaged versions λ_k^π and μ_{kl}^π .) However, modulation induces markedly different fluctuations *around* this fluid limit, i.e., in the diffusion scaling. We explore this in the next section.

4. CONVERGENCE RESULTS: A DIFFUSION LIMIT

Our main goal in this section is to establish a diffusion limit for the queue content process of the network under an appropriate scaling (which differs from the scaling imposed in the fluid limit result). This diffusion limit depends explicitly on scaling properties of the background process, thus giving insight into how the background process influences the behavior of the network on process level.

We will use several different tools to prove the diffusion limit. Typically, diffusion limits are proven relying on martingale methods and the Continuous Mapping Theorem (CMT); see, for instance, [20, 22, 23]. Unfortunately, the classical continuous maps do not apply to our network, due to the presence of the background process. We may get around this problem by constructing a more refined continuous map, but we opt for another solution. A crucial step in this solution consists of showing the asymptotic equivalence of the scaled network queue content process and a closely related stochastic process, which turns out to have the desired weak convergence properties. This approach combines queuing intuition with well-known continuous maps and convergence properties of stochastic integrals to achieve these results.

The continuous map that we will use is (a multidimensional version of) the well-known integral map presented in [20, Th. 4.1]. It is intimately related to standard infinite-server queues (cf. [20] and Lemma 3.2). For completeness, we present the result here. Its proof is a straightforward generalization of the proof of [20, Th. 4.1].

Lemma 4.1. *Let $h: \mathbb{R}^L \rightarrow \mathbb{R}^L$ be Lipschitz continuous, i.e.,*

$$\|h(a_1) - h(a_2)\| \leq c \|a_1 - a_2\|$$

for all $a_1, a_2 \in \mathbb{R}^L$ for some $c > 0$. Also assume that $h(0) = 0$. If $b \in \mathbb{R}^L$ and $x \in D([0, \infty); \mathbb{R}^L)$, then there exists a unique element $y \in D([0, \infty); \mathbb{R}^L)$ satisfying the integral equation

$$y(t) = b + x(t) + \int_0^t h(y(s)) ds. \quad (7)$$

Thus, the integral equation (7) defines a function $H: \mathbb{R}^L \times D([0, \infty); \mathbb{R}^L) \rightarrow D([0, \infty); \mathbb{R}^L)$. The function H is continuous whenever the space $D([0, \infty); \mathbb{R}^L)$ is equipped (in both the domain and the range) either with the uniform topology or with the weak J_1 topology.

The next corollary describes an example of a function H that is defined through the procedure in Lemma 4.1. The function H defined in this corollary is, in fact, the map that we will use to establish the diffusion limit. Moreover, H is clearly related to the fluid limit ρ (to see this, compare Eq. (6) and Eq. (8)).

Corollary 4.2. *Define the $L \times L$ matrix M via*

$$M_{kl} = \begin{cases} \mu_{lk}^\pi & \text{if } k \in \{1, \dots, L\}, l \in \mathcal{I}_k; \\ -\sum_{l \in \mathcal{I}_k} \mu_{kl}^\pi & \text{if } k \in \{1, \dots, L\}, l = k. \end{cases}$$

Let $h: \mathbb{R}^L \rightarrow \mathbb{R}^L$ be given by $h(a) = Ma$. Then h is Lipschitz continuous and Eq. (7) defines a function $H: \mathbb{R}^L \times D([0, \infty); \mathbb{R}^L) \rightarrow D([0, \infty); \mathbb{R}^L)$ that has the continuity properties described in Lemma 4.1. If $y = H(b, x)$, then

$$y_k(t) = b_k + x_k(t) + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}^\pi y_l(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}^\pi y_k(s) ds. \quad (8)$$

Throughout this section, the following conditions are in force, with

$$\beta = \max\{1/2, 1 - \alpha/2\}. \quad (9)$$

Assumption 4.3. *In the n -th system,*

- with $Q^n(0)$ denoting the vector with the initial number of jobs in queue k , we assume that

$$n^{1-\beta} \left(\frac{1}{n} Q^n(0) - \rho(0) \right) \rightarrow X(0) \quad (10)$$

in distribution, where $X(0)$ is some random variable;

- with λ_k^n the arrival rate to queue k , we assume that

$$\lim_{n \rightarrow \infty} n^{1-\beta} \left(\frac{1}{n} \lambda_k^n - \lambda_k \right) = \hat{\lambda}_k; \quad (11)$$

- with μ_{kl}^n the service speed for jobs to be routed from queue k to queue l , we assume that

$$\lim_{n \rightarrow \infty} n^{1-\beta} (\mu_{kl}^n - \mu_{kl}) = \hat{\mu}_{kl}; \quad (12)$$

- the background process, denoted by J^n , corresponds to a continuous-time Markov chain having irreducible generator matrix $n^\alpha \mathcal{Q}$ (for some $\alpha > 0$) and invariant distribution π .

Note that the assumption concerning the initial number of jobs is compatible with $\frac{1}{n} Q^n(0)$ converging in probability to $\rho(0)$, due to Slutsky's Lemma and $\beta - 1$ being strictly negative. In fact, for these reasons (10) implies that $\frac{1}{n} Q^n(0)$ converges in probability to $\rho(0)$.

The process of interest here is an appropriately centered and normalized version of Q^n , namely the stochastic process \hat{Q}^n , which is defined via

$$\hat{Q}_k^n(t) = n^{1-\beta} \left(\frac{1}{n} Q_k^n(t) - \rho_k(t) \right). \quad (13)$$

In words, \hat{Q}^n is the scaled and centered \mathbb{R}^L -valued process describing how the number of jobs in the network fluctuates around the fluid limit. As announced before, our objective is to show that \hat{Q}^n converges to a diffusion process. Anticipating the application of continuous-mapping methods, we represent \hat{Q}^n as

$$\hat{Q}_k^n(t) = \hat{Q}_k^n(0) + \hat{X}_k^n(t) + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}(J^n(s)) \hat{Q}_l^n(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}(J^n(s)) \hat{Q}_k^n(s) ds, \quad (14)$$

where

$$\hat{X}_k^n(t) = n^{1-\beta} \bar{X}_k^n(t) \quad (15)$$

and

$$\begin{aligned} \bar{X}_k^n(t) &= \bar{X}_{1,k}^n(t) + \bar{X}_{2,k}^n(t) + \bar{X}_{3,k}^n(t) + \sum_{l \in \mathcal{I}_k} (\bar{X}_{4,lk}^n(t) + \bar{X}_{5,lk}^n(t) + \bar{X}_{6,lk}^n(t)) \\ &\quad - \sum_{l \in \mathcal{I}_k} (\bar{X}_{4,kl}^n(t) + \bar{X}_{5,kl}^n(t) + \bar{X}_{6,kl}^n(t)), \end{aligned} \quad (16)$$

with the processes $\bar{X}_{1,k}^n$, $\bar{X}_{2,k}^n$, $\bar{X}_{3,k}^n$, $\bar{X}_{4,kl}^n$, $\bar{X}_{5,kl}^n$, and $\bar{X}_{6,kl}^n$ as defined in Eq. (22) (see Section 5.2).

Observe that the term $\mu_{kl}(J^n(s))$ in the integral $\int_0^t \mu_{kl}(J^n(s)) \widehat{Q}_l^n(s) ds$ above makes a direct application of Lemma 4.1 impossible. We deal with this problem by first showing that \widehat{Q}^n is asymptotically equivalent to the stochastic process \tilde{Q}^n , which is defined via

$$\tilde{Q}_k^n(t) = \widehat{Q}_k^n(0) + \widehat{X}_k^n(t) + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}^\pi \tilde{Q}_l^n(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}^\pi \tilde{Q}_k^n(s) ds \quad (17)$$

or, equivalently,

$$\tilde{Q}^n = H(\widehat{Q}^n(0), \widehat{X}^n), \quad (18)$$

with H denoting the continuous integral map from Example 4.2. Observe that (17) differs from (16) in that the modulated server speeds $\mu_{kl}(J^n(s))$ are replaced by their time-average counterparts. The equivalence of both systems is motivated by the following intuitive reasoning.

Our idea is to establish the diffusion result relying on classical continuous-mapping arguments: we would like to view \widehat{Q}^n as a continuous integral map of some input processes. If the server speeds were not modulated, it would be clear that the input is given by $\widehat{Q}^n(0)$ and \widehat{X}^n , and that \widehat{Q}^n is a continuous integral map of these input processes. Then everything would fall nicely within the scope of Lemma 4.1. Unfortunately, the server speeds do depend on the background process, so this approach does not work.

However, notice that the service requirements are not scaled, whereas the background process is scaled proportionally to n^α . Intuitively, this means that a service requirement remains fixed as n becomes large, while the background Markov chain is jumping faster and faster and approaches its equilibrium distribution. Thus, as n becomes large, a service requirement is expected to effectively experience the average or equilibrium server speed. These arguments suggest that, for large n , we may replace the modulated server speeds $\mu_{kl}(J^n(s))$ by the average server speeds μ_{kl}^π . More precisely, this means that the processes \widehat{Q}^n and \tilde{Q}^n should be asymptotically equivalent.

There are two major flaws in the reasoning above. First, \widehat{Q}^n is a scaled and centered queueing process rather than an actual queueing process; it is unclear what a ‘service requirement’ is in this context. Second, and more importantly, the term $\mu_{kl}(J^n(s))$ also appears in the input via the \widehat{X}^n (see Eq. (22) in Section 5.2). So why not take the average server speed there? The not so convincing answer is that in the input $\mu_{kl}(J^n(s))$ does not act as a server speed but is more like a measure of how the parameters deviate from their averaged counterparts.

Again, the intuitive reasoning above is not entirely convincing from a mathematical point of view, but at least it has given us a potentially useful idea. We formalize that idea in the following lemma.

Lemma 4.4. *The process \tilde{Q}^n is asymptotically equivalent to \widehat{Q}^n , meaning that $\widehat{Q}^n - \tilde{Q}^n$ converges uoc in probability to the zero process as $n \rightarrow \infty$.*

This lemma is useful because it has the following important corollary as an immediate consequence. Actually, the next result is the very reason for introducing \tilde{Q}^n : it claims that proving weak convergence of \widehat{Q}^n is equivalent to proving weak convergence of the simpler process \tilde{Q}^n , and that both processes must have the same limit.

Corollary 4.5. *The process \widehat{Q}^n converges weakly to X if and only if \widetilde{Q}^n converges weakly to X .*

Thus, to prove weak convergence of \widehat{Q}^n it suffices to prove weak convergence of \widetilde{Q}^n . However, because the process \widetilde{Q}^n was defined as a continuous function of some input process (cf. Eq. (18)), we only have to prove weak convergence of this input process. These observations are the basis of the proof of the following theorem, which is the main result of this section.

The statement of the theorem contains two equivalent systems of stochastic integral equations. In the first system, each term corresponds to the term in Eq. (14) or Eq. (16) from which it is derived. In the second system, we have reordered the terms for easier interpretation and to show how the form of the integral map H reappears in different parts of the limit.

Theorem 4.6. *Under Assumption 4.3, the processes \widetilde{Q}^n and \widehat{Q}^n both converge weakly to the unique solution X of the system of stochastic integral equations*

$$\begin{aligned}
X_k(t) = & X_k(0) + \mathbf{1}_{\{\alpha \geq 1\}} \sqrt{\lambda_k^\pi} B_{1,k}(t) + \widehat{\lambda}_k^\pi t + \mathbf{1}_{\{\alpha \leq 1\}} \lambda_k^\top B_b(t) \\
& + \sum_{l \in \mathcal{I}_k} \left(\mathbf{1}_{\{\alpha \geq 1\}} \int_0^t \sqrt{\mu_{lk}^\pi \rho_l(s)} dB_{lk}(s) + \int_0^t \widehat{\mu}_{lk}^\pi \rho_l(s) ds + \mathbf{1}_{\{\alpha \leq 1\}} \mu_{lk}^\top \int_0^t \rho_l(s) dB_b(s) \right) \\
& - \sum_{l \in \mathcal{I}_k} \left(\mathbf{1}_{\{\alpha \geq 1\}} \int_0^t \sqrt{\mu_{kl}^\pi \rho_k(s)} dB_{kl}(s) + \int_0^t \widehat{\mu}_{kl}^\pi \rho_k(s) ds + \mathbf{1}_{\{\alpha \leq 1\}} \mu_{kl}^\top \int_0^t \rho_k(s) dB_b(s) \right) \\
& + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}^\pi X_l(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}^\pi X_k(s) ds
\end{aligned} \tag{19}$$

or, equivalently,

$$\begin{aligned}
X_k(t) = & X_k(0) + \int_0^t \left(\widehat{\lambda}_k^\pi + \sum_{l \in \mathcal{I}_k} \widehat{\mu}_{lk}^\pi \rho_l(s) - \sum_{l \in \mathcal{I}_k} \widehat{\mu}_{kl}^\pi \rho_k(s) \right) ds \\
& + \mathbf{1}_{\{\alpha \leq 1\}} \int_0^t \left(\lambda_k^\top + \sum_{l \in \mathcal{I}_k} \mu_{lk}^\top \rho_l(s) - \sum_{l \in \mathcal{I}_k} \mu_{kl}^\top \rho_k(s) \right) dB_b(s) \\
& + \mathbf{1}_{\{\alpha \geq 1\}} \int_0^t \left(\sqrt{\lambda_k^\pi} dB_{1,k}(s) + \sum_{l \in \mathcal{I}_k} \sqrt{\mu_{lk}^\pi \rho_l(s)} dB_{lk}(s) - \sum_{l \in \mathcal{I}_k} \sqrt{\mu_{kl}^\pi \rho_k(s)} dB_{kl}(s) \right) \\
& + \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{lk}^\pi X_l(s) ds - \sum_{l \in \mathcal{I}_k} \int_0^t \mu_{kl}^\pi X_k(s) ds.
\end{aligned} \tag{20}$$

Here, the processes $B_{1,k}$ and B_{kl} are independent standard Brownian motions. The process B_b is an independent d -dimensional Brownian motion with $\langle B_b \rangle(t) = \Sigma t$, where $\Sigma := \text{diag}(\pi)D + D^\top \text{diag}(\pi)$ and the matrix D denotes the deviation matrix corresponding to the transition rate matrix \mathcal{Q} (as defined in Section 5.1).

5. PROOFS AND AUXILIARY RESULTS

In this section, we provide the proofs of the three main results, namely the fluid limit (Lemma 3.2), asymptotic equivalence (Lemma 4.4), and the diffusion limit (Theorem 4.6). Additionally, we present several auxiliary results which are used in those proofs. This section is organized in the following way.

First, we summarize some basic properties concerning the convergence of the state occupation measure of a Markov chain. These properties are essential ingredients of the upcoming proofs. We start with the proof of the fluid limit. This result is derived under slightly weaker conditions than the diffusion limit. Moreover, it is used in the other two proofs. Then, we proceed by establishing the asymptotic equivalence result in Lemma 4.4. Finally, relying on both the fluid limit and the asymptotic equivalence result, we give the proof of the diffusion limit.

5.1. State occupation measures of Markov chains. In the three upcoming proofs, the behavior of the state occupation measure of a scaled Markov chain plays a key role. In particular, we exploit weak convergence results for stochastic integrals with respect to the state occupation measure. To make this paper more or less self-contained, we collect the most important results concerning state occupation measures here. Their proofs can be found in [16].

As usual, J denotes a continuous-time Markov chain with state space $\{1, \dots, d\}$ for some $d \in \mathbb{N}$. It has $d \times d$ generator matrix \mathcal{Q} , which we assume to be irreducible. We denote the stationary distribution corresponding to \mathcal{Q} by π , i.e., π is the unique $d \times 1$ vector (with nonnegative entries summing to 1) that solves the equation $\pi^\top \mathcal{Q} = 0$. Additionally, D denotes the deviation matrix corresponding to \mathcal{Q} . The entries of D are given by

$$D_{ij} = \int_0^\infty (\mathbb{P}(J(s) = j \mid J(0) = i) - \pi_j) ds.$$

For some background on deviation matrices, see [7]. Define $\Sigma = \text{diag}(\pi)D + D^\top \text{diag}(\pi)$.

The state indicator function of J is the \mathbb{R}^d -valued function K defined via

$$K(i; t) = \mathbb{1}_{\{J(t)=i\}}$$

for $i \in \{1, \dots, d\}$ and $t \geq 0$. Then the state occupation measure corresponding to J is given by

$$\int_0^t K(s) ds.$$

Thus, the state indicator function K is used to keep track of the state of J at a given time, whereas the state occupation measure records how much time J has spent in each state during a time interval.

The following theorem concerns the behavior of the state occupation measure of a Markov chain when the generator matrix is scaled. Let \mathcal{Q} , π , D , and Σ be as defined above. Let J^n be a continuous-time Markov chain with generator matrix $n^\alpha \mathcal{Q}$ (for some $\alpha > 0$) and state indicator function K^n . Define the stochastic process G^n via

$$G^n(t) = \int_0^t (K^n(s) - \pi) ds. \tag{21}$$

Additionally, let Y^n denote the Dynkin martingale corresponding to K^n (cf. [1, Lem. 2.6.18]).

Theorem 5.1. *The process G^n converges to the zero process uoc in probability for $n \rightarrow \infty$. Additionally, $n^{\alpha/2} G^n$ converges weakly to a Brownian motion B_b having predictable quadratic variation process $\langle B_b \rangle(t) = \Sigma t$.*

In the proof of the diffusion limit, we also need weak convergence results of the form

$$H^n \cdot n^{\alpha/2} G^n \Rightarrow H \cdot G,$$

where $Y \cdot X$ denotes the Itô integral of Y with respect to X . This type of convergence generally does not hold, even when G is a Brownian motion, H is the zero process, and H^n converges to H uoc almost surely. However, under even stronger conditions this type of convergence does hold. We state such conditions in the next theorem.

Before we can state the theorem, we need to introduce some more notation (which by and large follows [15, p. 204]). Let X be a d -dimensional locally square-integrable martingale with respect to a filtration \mathbb{F} . For simplicity, we assume that

$$\langle X \rangle(t) = \int_0^t C(s) ds$$

for a predictable process C taking values in the set of all symmetric nonnegative $d \times d$ matrices. We denote by $L_{\text{loc}}^2(X)$ the set of predictable processes H taking values in $\mathbb{R}^{k \times d}$ such that the process

$$\int_0^t H(s)C(s)H(s)^\top ds$$

is locally integrable.

Theorem 5.2. *Consider the setting of Theorem 5.1. For fixed $m \in \mathbb{N}$, let $H^{1,n}, \dots, H^{m,n}$ be càdlàg processes in $L_{\text{loc}}^2(Y^n)$ and define $H_-^{j,n}$ via $H_-^{j,n}(t) = H^{j,n}(t-)$. Assume that each $H^{j,n}$ converges to a deterministic continuous function H^j uoc in probability and that $n^{-\alpha/2}H^{j,n}$ is a finite variation process whose total variation process converges to the zero process uoc in probability. Then $(H_-^{1,n} \cdot n^{\alpha/2}G^n, \dots, H_-^{m,n} \cdot n^{\alpha/2}G^n)$ converges weakly to $(H^1 \cdot B_b, \dots, H^m \cdot B_b)$.*

5.2. Proof of the fluid limit. To establish the fluid limit in Lemma 3.2, we follow a standard approach by combining the Functional Law of Large Numbers (FLLN) for Poisson processes with an application of Gronwall's Lemma. The fact that the parameters are modulated by a background process is a minor complication here, but we can deal with that by exploiting convergence properties of irreducible Markov chains (cf. Section 5.1).

Impose Assumption 3.1, as introduced in Section 3. Under these conditions, Lemma 3.2 claims uoc convergence of \bar{Q}_k^n to the zero process, where

$$\bar{Q}_k^n(t) = \frac{1}{n}Q_k^n(t) - \rho_k(t).$$

We prove this claim in the following way.

Using the expressions in Eq. (5) and Eq. (6), some simple algebra gives us

$$\begin{aligned} \bar{Q}_k^n(t) &= \bar{Q}_k^n(0) + \bar{X}_{1,k}^n(t) + \bar{X}_{2,k}^n(t) + \bar{X}_{3,k}^n(t) \\ &\quad + \sum_{l \in \mathcal{I}_k} \left(\bar{X}_{4,lk}^n(t) + \bar{X}_{5,lk}^n(t) + \int_0^t \mu_{lk}(J^n(s)) \bar{Q}_l^n(s) ds + \bar{X}_{6,lk}^n(t) \right) \\ &\quad - \sum_{l \in \mathcal{I}_k} \left(\bar{X}_{4,kl}^n(t) + \bar{X}_{5,kl}^n(t) + \int_0^t \mu_{kl}(J^n(s)) \bar{Q}_k^n(s) ds + \bar{X}_{6,kl}^n(t) \right), \end{aligned}$$

with

$$\begin{aligned} \bar{X}_{1,k}^n(t) &= \frac{1}{n}A_k \left(n \int_0^t \frac{1}{n} \lambda_k^n(J^n(s)) ds \right) - \int_0^t \frac{1}{n} \lambda_k^n(J^n(s)) ds, \\ \bar{X}_{2,k}^n(t) &= \int_0^t \frac{1}{n} \lambda_k^n(J^n(s)) ds - \int_0^t \lambda_k(J^n(s)) ds, \\ \bar{X}_{3,k}^n(t) &= \int_0^t \lambda_k(J^n(s)) ds - \int_0^t \lambda_k^\pi ds, \\ \bar{X}_{4,kl}^n(t) &= \frac{1}{n}S_{kl} \left(n \int_0^t \mu_{kl}^n(J^n(s)) \frac{1}{n} Q_k^n(s) ds \right) - \int_0^t \mu_{kl}^n(J^n(s)) \frac{1}{n} Q_k^n(s) ds, \\ \bar{X}_{5,kl}^n(t) &= \int_0^t (\mu_{kl}^n(J^n(s)) - \mu_{kl}(J^n(s))) \frac{1}{n} Q_k^n(s) ds, \\ \bar{X}_{6,kl}^n(t) &= \int_0^t (\mu_{kl}(J^n(s)) - \mu_{kl}^\pi) \rho_k(s) ds. \end{aligned} \tag{22}$$

Now suppose that all six types of the above processes converge to the zero process uoc in probability when $\frac{1}{n}Q^n(0)$ converges to $\rho(0)$ in probability. Then, for fixed $T \geq 0$, the random variables

$$\begin{aligned}\bar{Y}_1^n(T) &= \sum_{k=1}^L \left(\sup_{t \in [0, T]} |\bar{X}_{1,k}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{2,k}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{3,k}^n(t)| \right), \\ \bar{Y}_2^n(T) &= \sum_{k=1}^L \sum_{l \in \mathcal{I}_k} \left(\sup_{t \in [0, T]} |\bar{X}_{4,kl}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{5,kl}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{6,kl}^n(t)| \right),\end{aligned}$$

and

$$\bar{Y}_3^n(T) = \sum_{k=1}^L \sum_{l \in \mathcal{I}_k} \left(\sup_{t \in [0, T]} |\bar{X}_{4,kl}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{5,kl}^n(t)| + \sup_{t \in [0, T]} |\bar{X}_{6,kl}^n(t)| \right)$$

converge to 0 in the same way. Observe that

$$\begin{aligned}\sum_{k=1}^L |\bar{Q}_k^n(t)| &\leq \sum_{k=1}^L |\bar{Q}_k^n(0)| + \bar{Y}_1^n(T) + \bar{Y}_2^n(T) + \bar{Y}_3^n(T) \\ &\quad + \sum_{k=1}^L \sum_{l \in \mathcal{I}_k} \int_0^t \mu^* |\bar{Q}_l^n(s)| ds + \sum_{k=1}^L \sum_{l \in \mathcal{I}_k} \int_0^t \mu^* |\bar{Q}_k^n(s)| ds\end{aligned}$$

for all $t \in [0, T]$, where $\mu^* = \sum_{k=1}^L \sum_{l \in \mathcal{I}_k} \sum_{i=1}^d \mu_{kl}(i)$. It follows that

$$\sum_{k=1}^L |\bar{Q}_k^n(t)| \leq \sum_{k=1}^L |\bar{Q}_k^n(0)| + \bar{Y}_1^n(T) + \bar{Y}_2^n(T) + \bar{Y}_3^n(T) + \int_0^t 2(L-1)\mu^* \sum_{k=1}^L |\bar{Q}_k^n(s)| ds$$

for all $t \in [0, T]$. An application of Gronwall's Lemma (cf. [10, Th. A.5.1]) then leads to

$$\sum_{k=1}^L |\bar{Q}_k^n(t)| \leq \left(\sum_{k=1}^L |\bar{Q}_k^n(0)| + \bar{Y}_1^n(T) + \bar{Y}_2^n(T) + \bar{Y}_3^n(T) \right) e^{2(L-1)\mu^* T}$$

for all $t \in [0, T]$, so $\sup_{t \in [0, T]} \sum_{k=1}^L |\bar{Q}_k^n(t)|$ converges to 0 in probability under the assumptions imposed (in particular, recall that $\bar{Q}^n(0) \rightarrow 0$ in probability).

Consequently, it remains to show that $\bar{Y}_1^n(T)$, $\bar{Y}_2^n(T)$, and $\bar{Y}_3^n(T)$ converge to 0 in probability as soon as $\frac{1}{n}Q^n(0)$ converges to $\rho(0)$ in probability. To this end, we first define

$$\lambda^* = 1 + \sum_{k=1}^L \sum_{i=1}^d \lambda_k(i).$$

Then $\int_0^t \frac{1}{n} \lambda_k^n(J^n(s)) ds \leq 2\lambda^* t$ for all n large enough. Moreover, $\frac{1}{n} A_k(n, 2\lambda^* t)$ converges uoc to $2\lambda^* t$ almost surely by the FLLN for Poisson processes (cf. [6, Th. 5.5.10]), implying that $\sup_{t \in [0, T]} |\bar{X}_{1,k}^n(t)|$ converges to 0 almost surely.

Next, observe that $\sup_{t \in [0, T]} |\bar{X}_{2,k}^n(t)|$ converges to 0 almost surely, due to the convergence of $\frac{1}{n} \lambda_k^n$ to λ_k . Also observe that both $\sup_{t \in [0, T]} |\bar{X}_{3,k}^n(t)|$ and $\sup_{t \in [0, T]} |\bar{X}_{6,kl}^n(t)|$ converge to 0 in probability by Theorem 5.1 and Theorem 5.2, respectively.

To analyze $\bar{X}_{4,kl}^n(t)$ and $\bar{X}_{5,kl}^n(t)$, we need to know more about the behavior of $Q_k^n(s)$. Clearly, the queueing dynamics are such that

$$\left| \frac{1}{n} Q_k^n(s) \right| \leq \sum_{m \in \mathcal{I}} \left(\frac{1}{n} Q_m^n(0) + \frac{1}{n} A_m(n, \lambda^* T) \right)$$

for all $s \in [0, T]$. The sum on the right-hand side converges to $\sum_{m \in \mathcal{I}} (\rho_m(0) + \lambda^* T)$ in probability, due to Slutsky's Lemma (cf. [21, Lem. 2.8]). Another application of Slutsky's Lemma shows that

$$|\mu_{kl}^n(J^n(s)) - \mu_{kl}(J^n(s))| \sum_{l=1}^L \left(\frac{1}{n} Q_l^n(0) + \frac{1}{n} A_l(n\lambda^* T) \right)$$

converges to 0 in probability. Then $\sup_{t \in [0, T]} |\bar{X}_{5,kl}^n(t)|$ must converge to 0 in probability, too. Regarding $\bar{X}_{4,kl}^n(t)$, we observe that, for fixed $\epsilon > 0$, we have the very crude inequality

$$\int_0^t \mu_{kl}^n(J^n(s)) \frac{1}{n} Q_k^n(s) ds \leq \mu^* \sum_{m \in \mathcal{I}} (\rho_m(0) + \lambda^* T) T + \epsilon$$

on a set with probability at least $1 - \epsilon$ for all n large enough. Combined with the FLLN for the standard Poisson process S_{kl} , this gives us the convergence of $\sup_{t \in [0, T]} |\bar{X}_{4,kl}^n(t)|$ to 0 in probability.

Using Slutsky's Lemma once again, we conclude that $\bar{Y}_1^n(T)$, $\bar{Y}_2^n(T)$, and $\bar{Y}_3^n(T)$ converge to 0 in probability. This completes the proof.

5.3. Proof of asymptotic equivalence. Impose Assumption 4.3, as stated in Section 4. Our goal is to prove the claim in Lemma 4.4, namely that \tilde{Q}^n is asymptotically equivalent to \hat{Q}^n . Recall that \hat{Q}^n represents the scaled and centered queue content process and that \tilde{Q}^n is a similar but simpler stochastic process defined using a continuous map. (Note that, in general, \tilde{Q}^n is not a scaled and centered queue content process.)

For completeness, we first show that \tilde{Q}^n is a well-defined stochastic process. We already know that $\hat{Q}^n(0)$ and \hat{X}^n are well-defined random elements. As observed in Example 4.2, the integral map H that defines \tilde{Q}^n satisfies the conditions of Lemma 4.1, so \tilde{Q}^n is a continuous function (in the Skorokhod J_1 topology) of a stochastic process. Then \tilde{Q}^n is a well-defined stochastic process.

We turn to proving asymptotic equivalence of \hat{Q}^n and \tilde{Q}^n . The difference between these two processes can be written as

$$\begin{aligned} & \hat{Q}_k^n(t) - \tilde{Q}_k^n(t) \\ &= \sum_{l \in \mathcal{I}_k} \left(\int_0^t \mu_{lk}(J^n(s)) \hat{Q}_l^n(s) ds - \int_0^t \mu_{lk}^\pi \tilde{Q}_l^n(s) ds \right) \\ & \quad - \sum_{l \in \mathcal{I}_k} \left(\int_0^t \mu_{kl}(J^n(s)) \hat{Q}_k^n(s) ds - \int_0^t \mu_{kl}^\pi \tilde{Q}_k^n(s) ds \right) \\ &= \sum_{l \in \mathcal{I}_k} \left(\int_0^t (\mu_{lk}(J^n(s)) - \mu_{lk}^\pi) \hat{Q}_l^n(s) ds - \int_0^t \mu_{lk}^\pi [\tilde{Q}_l^n(s) - \hat{Q}_l^n(s)] ds \right) \\ & \quad - \sum_{l \in \mathcal{I}_k} \left(\int_0^t (\mu_{kl}(J^n(s)) - \mu_{kl}^\pi) \hat{Q}_k^n(s) ds - \int_0^t \mu_{kl}^\pi [\tilde{Q}_k^n(s) - \hat{Q}_k^n(s)] ds \right), \end{aligned}$$

so this difference is, in fact, given by the continuous integral map H of some input process. Consequently, if this input process converges to the zero process uoc in probability, then the difference between \hat{Q}^n and \tilde{Q}^n must converge to the zero process uoc in probability, too.

This input process is essentially a sum of other input processes of the form

$$\int_0^t (\mu_{kl}(J^n(s)) - \mu_{kl}^\pi) \hat{Q}_k^n(s) ds. \quad (23)$$

Thus, to prove the lemma, it suffices to establish the convergence of each of these terms to the zero process uoc in probability. We will do this in the remainder of this proof.

Recall from Section 5.1 that the state indicator function of J^n is denoted by K^n and that the Dynkin martingale associated with J^n is denoted by Y^n . Now rewrite the integral in Eq. (23) as

$$\int_0^t n^{-\alpha/2} \widehat{Q}_k^n(s) \mu_{kl}^\top n^{\alpha/2} (K^n(s) - \pi) ds = \int_0^t U_{kl}^n(s) n^{\alpha/2} (K^n(s) - \pi) ds, \quad (24)$$

where

$$U_{kl}^n(t) = n^{-\alpha/2} \widehat{Q}_k^n(t-) \mu_{kl}^\top.$$

Convergence of U_{kl}^n to the zero process implies, under certain assumptions, convergence of the integral (24) to the zero process; this is described Theorem 5.2.

To be able to apply Theorem 5.2, we have to check several properties. We will break this up in four steps and then apply Theorem 5.2 in the fifth and final step.

Step 1. We would like to verify (in this step and the next) that U_{kl}^n is a process in $L_{\text{loc}}^2(Y^n)$. First of all, we should have a filtration with respect to which U_{kl}^n is predictable and Y^n is a martingale. For this filtration, say \mathbb{G} , we can just take the natural filtration of J^n and add to this the entire filtration of all Poisson processes and $Q^n(0)$, which are independent of J^n . Then Y^n is still a (locally square-integrable) martingale with respect to this filtration, due to the independence. Moreover, the left-continuous process U_{kl}^n is adapted to \mathbb{G} and thus predictable. This settles the filtration and measurability issues of Theorem 5.2.

Step 2. From the previous step, we know that U_{kl}^n has the right measurability properties. To verify that U_{kl}^n is a process in $L_{\text{loc}}^2(Y^n)$, we also have to check that the process

$$\int_0^t U_{kl}^n(s) \left[\text{diag}(\mathcal{Q}^\top K(s)) - \mathcal{Q}^\top \text{diag}(K(s)) - \text{diag}(K(s)) \mathcal{Q} \right] U_{kl}^n(s)^\top ds$$

is locally integrable. This is rather obvious, as

$$\begin{aligned} \left| \widehat{Q}_k^n(t) \right| &\leq n^{1-\beta} \frac{1}{n} Q_k^n(t) + n^{1-\beta} \rho_k(t) \\ &\leq n^{1-\beta} \rho_k(t) + \sum_{l=1}^L A_l \left(\int_0^t \lambda_l^n(J^n(s)) ds \right) \\ &\leq n^{1-\beta} \rho_k(t) + \sum_{l=1}^L A_l(\lambda^* t), \end{aligned}$$

where we now define λ^* as $\sum_{k=1}^L \sum_{i=1}^d \lambda_k(i)$. From this, we conclude that U_{kl}^n is a process in $L_{\text{loc}}^2(Y^n)$.

Step 3. Perhaps the most important step is to check that U_{kl}^n converges to the zero process uoc in probability. From the definition of \widehat{Q}_k^n in Eq. (14), it follows that

$$U_{kl}^n(t) = n^{1-\beta-\alpha/2} \left(\frac{1}{n} Q_k^n(t-) - \rho_k(t-) \right) \mu_{kl}^\top. \quad (25)$$

By the fluid limit in Lemma 3.2, the process $\frac{1}{n} Q_k^n - \rho_k$ converges to the zero process uoc in probability. Because $1 - \beta - \alpha/2 = \min\{0, (1 - \alpha)/2\}$, we know that U_{kl}^n converges to the zero process in the same way.

Step 4. Now we only have to verify that $n^{-\alpha/2} U_{kl}^n$ is a finite variation process whose total variation process converges to the zero process uoc in probability. From Eq. (25), we obtain

$$\begin{aligned} n^{-\alpha/2} U_{kl}^n(t) &= n^{-\alpha/2} n^{1-\beta-\alpha/2} \left(\frac{1}{n} Q_k^n(t-) - \rho_k(t-) \right) \mu_{kl}^\top \\ &= n^{1-\beta-\alpha} \left(\frac{1}{n} Q_k^n(t-) - \rho_k(t-) \right) \mu_{kl}^\top. \end{aligned} \quad (26)$$

Clearly, it suffices to show that $n^{1-\beta-\alpha}(\frac{1}{n}Q_k^n(t-) - \rho_k(t-))$ is a finite variation process whose total variation process converges to the zero process uoc in probability.

To this end, observe that ρ_k is continuously differentiable, so it is of bounded variation. Moreover, ρ_k does not depend on n and we have

$$1 - \beta - \alpha = \min\{-\alpha/2, (1 - 2\alpha)/2\} < 0, \quad (27)$$

so the total variation process of $n^{1-\beta-\alpha}\rho_k$ converges to the zero process uoc. Consequently, it suffices to show that $n^{1-\beta-\alpha}\frac{1}{n}Q_k^n(t-)$ is a finite variation process whose total variation process converges to the zero process uoc in probability.

Now note that $\frac{1}{n}Q_k^n(t-)$ can be written as the difference of nondecreasing processes (cf. Eq. (5)). This means that $n^{1-\beta-\alpha}\frac{1}{n}Q_k^n(t-)$ is a finite variation process. It also follows from Eq. (5) that the total variation process of $n^{1-\beta-\alpha}\frac{1}{n}Q_k^n(t-)$ is given by

$$\begin{aligned} & n^{1-\beta-\alpha}\frac{1}{n}Q_k^n(0) + n^{1-\beta-\alpha}\frac{1}{n}A_k\left(n\int_0^t\frac{1}{n}\lambda_k^n(J^n(s))ds\right) \\ & + n^{1-\beta-\alpha}\sum_{l\in\mathcal{I}_k}\frac{1}{n}S_{lk}\left(n\int_0^t\mu_{lk}^n(J^n(s))\frac{1}{n}Q_l^n(s)ds\right) \\ & + n^{1-\beta-\alpha}\sum_{l\in\mathcal{I}_k}\frac{1}{n}S_{kl}\left(n\int_0^t\mu_{kl}^n(J^n(s))\frac{1}{n}Q_k^n(s)ds\right). \end{aligned}$$

As a result of Lemma 3.2 and the fact that $1 - \beta - \alpha = \min\{-\alpha/2, (1 - 2\alpha)/2\}$, we obtain the convergence of this expression to the zero process uoc in probability.

As argued before, this implies that $n^{-\alpha/2}U_{kl}^n$ is a finite variation process whose total variation process converges to the zero process uoc in probability.

Step 5. Recall that we aim to show convergence of the integral in Eq. (24) to the zero process via an application of Theorem 5.2. In the previous four steps, we have verified the assumptions of Theorem 5.2. Moreover, we have shown that U_{kl}^n converges to the zero process uoc in probability, so Theorem 5.2 implies that the integral in Eq. (24) converges to the zero process uoc in probability, as required.

5.4. Proof of the diffusion limit. Before we turn to the proof of the diffusion limit result of Theorem 4.6, let us first comment on notation. We will be mainly concerned with joint weak convergence of certain stochastic processes. Because this involves many different processes, we will not explicitly order them in a vector, so as to minimize potential confusion. As an example, we will just say that Y_k and Y_{kl} converge jointly, where it is tacitly understood that $k \in \{1, \dots, L\}$ and $l \in \mathcal{I}_k$ (the ranges of the indices should be clear from the context), and that all these processes are ordered in some vector.

Impose Assumption 4.3, as was stated in Section 4. We would like to establish the diffusion limit in Theorem 4.6 under these conditions. To this end, it suffices to show that \tilde{Q}^n converges weakly to the process X given in Theorem 4.6, due to Lemma 4.4. However, because \tilde{Q}^n is defined as a continuous function of certain input processes $\hat{Q}^n(0)$ and \hat{X}^n , the CMT implies that we only have to prove joint weak convergence of these input processes. Consequently, the remainder of this proof will focus on establishing this joint weak convergence.

Following the notation from Section 5.1, we denote by K^n the state indicator function corresponding to the background process J^n . Recall that K^n is a $\{0, 1\}^d$ -valued process that is given by $K^n(i; t) =$

$\mathbb{1}_{\{J^n(s)=i\}}$ for $i \in \{1, \dots, d\}$. In addition, we define the continuous semimartingale G^n via

$$G^n(t) = \int_0^t (K^n(s) - \pi) ds.$$

As argued before, we only have to prove weak convergence of the input processes of \tilde{Q}^n , i.e., we have to prove joint weak convergence of $\hat{Q}^n(0)$ and \hat{X}^n . We will show that weak convergence of \hat{X}^n follows (after an application of the CMT) from the joint weak convergence of other, more basic processes. Joint convergence of $\hat{Q}^n(0)$ and these basic processes will be straightforward (as they are independent of $\hat{Q}^n(0)$ or converge to a deterministic limit), so we will focus on the convergence of \hat{X}^n .

Recall from Eq. (15) and Eq. (16) that \hat{X}^n is built up from six different types of one-dimensional processes, namely

$$n^{1-\beta} \bar{X}_{1,k}^n, \quad n^{1-\beta} \bar{X}_{2,k}^n, \quad n^{1-\beta} \bar{X}_{3,k}^n, \quad n^{1-\beta} \bar{X}_{4,kl}^n, \quad n^{1-\beta} \bar{X}_{5,kl}^n, \quad \text{and} \quad n^{1-\beta} \bar{X}_{6,kl}^n. \quad (28)$$

Clearly, joint weak convergence of all these processes implies weak convergence of \hat{X}^n , as a result of the CMT.

To establish joint weak convergence of these processes, we will exploit the fact that some of them are independent or converge to deterministic functions. For those processes that are independent or converge to a deterministic function, joint convergence follows automatically (cf. [23, Th. 11.4.4] and [23, Th. 11.4.5]). We will prove joint weak convergence of the six processes in (28) in four steps. We will then summarize these steps in a fifth and final step.

Step 1. Some of the six types of processes (namely the the processes $n^{1-\beta} \bar{X}_{2,k}^n$ and $n^{1-\beta} \bar{X}_{5,kl}^n$) converge in probability to deterministic functions. Indeed,

$$n^{1-\beta} \bar{X}_{2,k}^n(t) = n^{1-\beta} \left(\frac{1}{n} \lambda_k^n - \lambda_k \right)^\top \int_0^t K^n(s) ds,$$

which converges uoc in probability to the deterministic process

$$\hat{\lambda}_k^\top \int_0^t \pi ds = \int_0^t \hat{\lambda}_k^\top \pi ds$$

by Eq. (11) and Theorem 5.1. The process $n^{1-\beta} \bar{X}_{5,kl}^n$ satisfies

$$\begin{aligned} n^{1-\beta} \bar{X}_{5,kl}^n(t) &= n^{1-\beta} \left(\frac{1}{n} \mu_{kl}^n - \mu_{kl} \right)^\top \int_0^t \frac{1}{n} Q_k^n(s) K^n(s) ds \\ &= n^{1-\beta} \left(\frac{1}{n} \mu_{kl}^n - \mu_{kl} \right)^\top \int_0^t \left(\frac{1}{n} Q_k^n(s) - \rho_k(s) \right) K^n(s) ds \\ &\quad + n^{1-\beta} \left(\frac{1}{n} \mu_{kl}^n - \mu_{kl} \right)^\top \int_0^t \rho_k(s) K^n(s) ds. \end{aligned}$$

Consider the last two terms above. The first of these converges uoc in probability to the zero process. This follows immediately from the fluid limit in Lemma 3.2 and the fact that K^n is bounded by 1. The convergence of the second term, on the other hand, follows from similar arguments as the convergence of $n^{1-\beta} \bar{X}_{2,k}^n$ by combining Eq. (12) and Theorem 5.1. In this case, the limiting process is given by

$$\int_0^t \mu_{kl}^\pi \rho_k(s) ds.$$

Consequently, $n^{1-\beta} \bar{X}_{5,kl}^n$ converges uoc in probability to this process, too.

Step 2. We have reduced our task to proving joint convergence of $n^{1-\beta}\bar{X}_{1,k}^n$, $n^{1-\beta}\bar{X}_{3,k}^n$, $n^{1-\beta}\bar{X}_{4,kl}^n$, and $n^{1-\beta}\bar{X}_{6,kl}^n$. A closer inspection of these processes reveals that each of them comes in one of two flavors. On the one hand, $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$ can be written as a stochastic integral with respect to the semimartingale G^n . On the other hand, $n^{1-\beta}\bar{X}_{1,k}^n$ and $n^{1-\beta}\bar{X}_{4,kl}^n$ can be viewed as a random time change of the Poisson martingales

$$n^{1-\beta}\bar{A}_k^n(t) = n^{1-\beta}\left(\frac{1}{n}A_k(nt) - t\right), \quad n^{1-\beta}\bar{S}_{kl}^n(t) = n^{1-\beta}\left(\frac{1}{n}S_{kl}(nt) - t\right),$$

respectively. Indeed, defining the random times

$$\tau_{1,k}^n(t) = \int_0^t \frac{1}{n}\lambda_k^n(J^n(s)) ds, \quad \tau_{4,kl}^n(t) = \int_0^t \mu_{kl}^n(J^n(s)) \frac{1}{n}Q_k^n(s) ds,$$

we see that $n^{1-\beta}\bar{X}_{1,k}^n = n^{1-\beta}\bar{A}_k^n \circ \tau_{1,k}^n$ and $n^{1-\beta}\bar{X}_{4,kl}^n = n^{1-\beta}\bar{S}_{kl}^n \circ \tau_{4,kl}^n$.

The idea is to prove joint weak convergence of $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$, together with $n^{1-\beta}\bar{A}_k^n$, $n^{1-\beta}\bar{S}_{kl}^n$, and the random times $\tau_{1,k}^n$ and $\tau_{4,kl}^n$. Say that they have limits $X_{3,k}$, $X_{6,kl}$, B_k , B_{kl} , $\tau_{1,k}$, and $\tau_{4,kl}$, respectively. Then, under mild conditions, [23, Th. 13.2.2] guarantees the joint weak convergence of $n^{1-\beta}\bar{X}_{3,k}^n$, $n^{1-\beta}\bar{X}_{6,kl}^n$, $n^{1-\beta}\bar{A}_k^n \circ \tau_{1,k}^n$, and $n^{1-\beta}\bar{S}_{kl}^n \circ \tau_{4,kl}^n$. Additionally, [23, Th. 13.2.2] identifies the corresponding limits as $X_{3,k}$, $X_{6,kl}$, $B_k \circ \tau_{1,k}$, and $B_{kl} \circ \tau_{4,kl}$. The conditions of this theorem will be clearly met in our case, as all limits will turn out to be continuous.

Step 3. Thus, before applying [23, Th. 13.2.2], we have to prove joint weak convergence of $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$, together with $n^{1-\beta}\bar{A}_k^n$, $n^{1-\beta}\bar{S}_{kl}^n$, and the random times $\tau_{1,k}^n$ and $\tau_{4,kl}^n$. To prove this joint weak convergence, observe that the processes $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$ are independent from the Poisson martingales, because the background process is independent from these martingales by assumption. Consequently, $n^{1-\beta}\bar{X}_{3,k}^n$, $n^{1-\beta}\bar{X}_{6,kl}^n$, and the Poisson martingales converge jointly as soon as $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$ converge jointly and the Poisson martingales converge jointly.

However, in general $n^{1-\beta}\bar{X}_{3,k}^n$, $n^{1-\beta}\bar{X}_{6,kl}^n$, and the Poisson martingales are not independent from $\tau_{1,k}^n$ and $\tau_{4,kl}^n$. Nevertheless, dealing with the random time changes is relatively straightforward: they converge uoc in probability to deterministic functions, which is a sufficient condition for joint weak convergence. Note that, similar to $n^{1-\beta}\bar{X}_{2,k}^n$ and $n^{1-\beta}\bar{X}_{5,kl}^n$, the process $\tau_{1,k}^n$ converges to the continuous function $\tau_{1,k}(t) = \int_0^t \lambda_k^\pi ds$ and the process $\tau_{4,kl}^n$ converges to the continuous function $\tau_{4,kl}(t) = \int_0^t \mu_{kl}^\pi \rho_k(s) ds$.

Now consider the Poisson martingales and recall that they are all independent. For ease of exposition, take $n^{1-\beta}\bar{A}_k^n = n^{1/2-\beta}\sqrt{n}\bar{A}_k^n$. We know from the MCLT that $\sqrt{n}\bar{A}_k^n$ converges weakly to a standard Brownian motion B_k (cf. [20, Th. 4.2]). Since $1/2-\beta = \min\{0, (\alpha-1)/2\}$, it follows that $n^{1-\beta}\bar{A}_k^n$ converges weakly to $\mathbf{1}_{\{\alpha \geq 1\}}B_k$. Similarly, $n^{1-\beta}\bar{S}_{kl}^n$ converges weakly to $\mathbf{1}_{\{\alpha \geq 1\}}B_{kl}$, where B_{kl} is a standard Brownian motion. Because of the Poisson martingales being independent, joint convergence holds as well and the Brownian motions B_k and B_{kl} are independent.

Step 4. It remains to show joint weak convergence of the processes $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$. We observed earlier that these processes are stochastic integrals with respect to the continuous semimartingale G^n . More concretely, we have

$$n^{1-\beta}\bar{X}_{3,k}^n(t) = \int_0^t n^{1-\beta-\alpha/2}\lambda_k^\top dn^{\alpha/2}G^n(s)$$

and

$$n^{1-\beta}\bar{X}_{6,kl}^n(t) = \int_0^t n^{1-\beta-\alpha/2}\rho_k(s-)\mu_{kl}^\top dn^{\alpha/2}G^n(s).$$

The form of these integrals suggests that we may apply Theorem 5.2.

Let us check that the conditions of Theorem 5.2 are satisfied here. First, recall the definition of β in Eq. (9). In particular, it holds that $1 - \beta - \alpha/2 = 0$ for $\alpha \leq 1$, $1 - \beta - \alpha/2 < 0$ for $\alpha > 1$, and $1 - \beta - \alpha < 0$. Next, observe that $n^{1-\beta-\alpha/2}\lambda_k^\top$ converges to $\mathbb{1}_{\{\alpha \leq 1\}}\lambda_k^\top$ uoc. In addition, $n^{1-\beta-\alpha/2}\rho_k\mu_{kl}^\top$ converges to $\mathbb{1}_{\{\alpha \leq 1\}}\rho_k\mu_{kl}^\top$ uoc. Third, notice that ρ_k is a function of bounded variation and that ρ_k does not depend on n . Consequently, both $n^{-\alpha/2}n^{1-\beta-\alpha/2}\lambda_k^\top$ and $n^{-\alpha/2}n^{1-\beta-\alpha/2}\rho_k\mu_{kl}^\top$ are finite-variation processes whose total variation processes converge to the zero process uoc.

Combining all these observations, we see that the conditions of Theorem 5.2 are satisfied. Then Theorem 5.2 implies the joint weak convergence of the stochastic integrals $n^{1-\beta}\bar{X}_{3,k}^n$ and $n^{1-\beta}\bar{X}_{6,kl}^n$ to

$$\int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\lambda_k^\top dB_b(s) \quad \text{and} \quad \int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\rho_k\mu_{kl}^\top dB_b(s),$$

respectively.

Step 5. Summarizing, we have shown joint weak convergence of the eight processes

$$n^{1-\beta}\bar{X}_{2,k}^n, \quad n^{1-\beta}\bar{X}_{5,kl}^n, \quad n^{1-\beta}\bar{A}_k^n, \quad n^{1-\beta}\bar{S}_{kl}^n, \quad \tau_{1,k}^n, \quad \tau_{4,kl}^n, \quad n^{1-\beta}\bar{X}_{3,k}^n, \quad \text{and} \quad n^{1-\beta}\bar{X}_{6,kl}^n$$

to, respectively,

$$\begin{aligned} & \int_0^t \widehat{\lambda}_k^\pi ds, \quad \int_0^t \mu_{kl}^\pi \rho_k(s) ds, \quad \mathbb{1}_{\{\alpha \geq 1\}}B_k, \quad \mathbb{1}_{\{\alpha \geq 1\}}B_{kl}, \quad \tau_{1,k}, \quad \tau_{4,kl}, \\ & \int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\lambda_k^\top dB_b(s), \quad \text{and} \quad \int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\rho_k\mu_{kl}^\top dB_b(s). \end{aligned}$$

This implies joint weak convergence of the six processes

$$n^{1-\beta}\bar{X}_{2,k}^n, \quad n^{1-\beta}\bar{X}_{5,kl}^n, \quad n^{1-\beta}\bar{X}_{1,k}^n, \quad n^{1-\beta}\bar{X}_{4,kl}^n, \quad n^{1-\beta}\bar{X}_{3,k}^n, \quad \text{and} \quad n^{1-\beta}\bar{X}_{6,kl}^n$$

to, respectively,

$$\begin{aligned} & \int_0^t \widehat{\lambda}_k^\pi ds, \quad \int_0^t \mu_{kl}^\pi \rho_k(s) ds, \quad \mathbb{1}_{\{\alpha \geq 1\}}B_k \circ \tau_{1,k}, \quad \mathbb{1}_{\{\alpha \geq 1\}}B_{kl} \circ \tau_{4,kl}, \\ & \int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\lambda_k^\top dB_b(s), \quad \text{and} \quad \int_0^t \mathbb{1}_{\{\alpha \leq 1\}}\rho_k\mu_{kl}^\top dB_b(s). \end{aligned}$$

Now applying the CMT, and using that $\mathbb{1}_{\{\alpha \geq 1\}}B_k \circ \tau_{1,k}$ and $\mathbb{1}_{\{\alpha \geq 1\}}B_{kl} \circ \tau_{4,kl}$ have the same distributional properties as the integrals

$$\mathbb{1}_{\{\alpha \geq 1\}} \int_0^t \sqrt{\lambda_k^\pi} dB_{1,k}(s) \quad \text{and} \quad \mathbb{1}_{\{\alpha \geq 1\}} \int_0^t \sqrt{\mu_{kl}^\pi \rho_k(s)} dB_{kl}(s),$$

the statement of Theorem 4.6 follows immediately.

REFERENCES

- [1] Lakhdar Aggoun and Robert Elliott. *Measure Theory and Filtering*. Cambridge University Press, New York, 2004.
- [2] D. Anderson, J. Blom, M. Mandjes, H. Thorsdottir, and K. De Turck. A functional central limit theorem for a Markov-modulated infinite-server queue. *Methodology and Computing in Applied Probability*, 18(1):153–168, 2016.
- [3] A. Bassamboo, R.S. Randhawa, and A. Zeevi. Capacity sizing under parameter uncertainty: safety staffing principles revisited. *Management Science*, 56, 2010.
- [4] M. Baykal-Gursoy and W. Xiao. Stochastic decomposition in M/M/ ∞ queues with Markov modulated service rates. *Queueing Systems*, 48(1):75–88, 2004.
- [5] J. Blom, K. De Turck, and M. Mandjes. Functional central limit theorems for Markov-modulated infinite-server systems. *Mathematical Methods of Operations Research*, 83(3):351–372, 2016.
- [6] Hong Chen and David D. Yao. *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*. Springer-Verlag, New York, 2001.

- [7] Pauline Coolen-Schrijner and Erik A. van Doorn. The deviation matrix of a continuous-time Markov chain. *Probability in the Engineering and Informational Sciences*, 16(3):351–366, 2002.
- [8] B. D’Auria. M/M/∞ queues in semi-Markovian random environment. *Queueing Systems*, 58(3):221–237, 2008.
- [9] S. G. Eick, W. A. Massey, and W. Whitt. The physics of the $M_t/G/\infty$ queue. *Operations Research*, 41:731–742, 1993.
- [10] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.
- [11] G. Falin. The M/M/∞ queue in a random environment. *Queueing Systems*, 58(1):65–76, 2008.
- [12] Brian H. Fralix and Ivo J. B. F. Adan. An infinite-server queue influenced by a semi-Markovian environment. *Queueing Systems*, 61(1):65–84, 2009.
- [13] J. Michael Harrison and Austin J. Lemoine. A note on networks of infinite-server queues. *Journal of Applied Probability*, 18(1):561–567, 1981.
- [14] Mariska Heemskerk, Johan van Leeuwen, and Michel Mandjes. Scaling limits for infinite-server systems in a random environment. *Stochastic Systems*, 7(1):1–31, 2017.
- [15] Jean Jacod and Albert N. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer-Verlag, Berlin Heidelberg, second edition, 2003.
- [16] H. M. Jansen. Weak convergence of stochastic integrals with respect to the state occupation measure of a Markov chain. arXiv:1712.03396v1, 2017.
- [17] S. Kim and W. Whitt. Are call center and hospital arrivals well modeled by nonhomogeneous Poisson processes? *Manufacturing & Service Operations Management*, 16:464–480, 2014.
- [18] D. T. Koops, O. J. Boxma, and M. R. H. Mandjes. Networks of $\cdot/G/\infty$ queues with shot-noise-driven arrival intensities. *Queueing Systems*, 86(3-4):301–325, 2017.
- [19] C. A. O’Cinneide and P. Purdue. The M/M/∞ queue in a random environment. *Journal of Applied Probability*, 23(1):175–184, 1986.
- [20] Guodong Pang, Rishi Talreja, and Ward Whitt. Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probability Surveys*, 4(1):193–267, 2007.
- [21] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, New York, 1998.
- [22] Ward Whitt. Some useful functions for functional limit theorems. *Mathematics of Operations Research*, 5(1):67–85, 1980.
- [23] Ward Whitt. *Stochastic-Process Limits: an Introduction to Stochastic-Process Limits and Their Application to Queues*. Springer, New York, 2002.